

Intrusion Detection Terminology (Part Two)

Andy Cuff 2003-09-24

The [first part](#) of this series discussed the concept of Alerts, Consoles, False Negatives, and many other terms that are important for Intrusion Detection Systems (IDS). This second and final terminology article will continue in the same vein, starting with an explanation of the many different types of IDSs that exist today.

IDS Categories

Although we tend to talk about IDSs as though they are just one thing, there are actually many different types of IDS. The following is a list of the various types of IDS, and a brief explanation of what differentiates them.

Application IDS

Application IDSs are aware of the intrusion signatures for specific applications, usually the more vulnerable applications such as Web servers, databases etc. However, many of the host-based IDSs that ordinarily look at operating systems are becoming more application-aware. One example of an application-specific IDS is [Entercept Web Server Edition](#).

Consoles -- see [Consoles](#) in Part One of this series.

File Integrity Checkers

When a system is compromised by an attacker, it will often alter certain key files to provide continued access and prevent detection. By applying a message digest (cryptographic hash) to key files, the files can be checked periodically to see if they have been altered, thus providing a degree of assurance. Upon detecting such a change, the file integrity checker will trigger an alert. The same process can be employed by a system administrator after being successfully attacked, allowing him/her to ascertain the extent to which the system has been compromised. Previously, file integrity checkers detected intrusions long after the event; however, more products have recently been emerging that check files as they are accessed, thereby introducing a near realtime IDS element.

Examples include [Tripwire](#) and [Intact](#).

Honeypots -- see [Honeypot](#) in Part One of this series

Host-based IDS

This kind of IDS monitors sys/event logs from multiple sources for suspicious activity. Host-based IDSs (also known as simply **Host IDS**) are best placed to detect computer misuse from trusted insiders and those who have infiltrated your network evading traditional methods of detection. What I've just described is really an event log viewer (with attitude). A true host IDS will apply some signature analysis across multiple events/logs and/or time. Many will also incorporate heuristics into the product. Some will introduce an added benefit: because they operate at near real time, system faults are often detected quickly, which makes them popular with techies as well as security personnel. The term host-based IDS has been applied to any kind of IDS sitting on a workstation/server. Vendors have tried this for various products from Network Node IDS to File Integrity Checkers, and while I can understand their logic, this can be misleading to the buyer.

Hybrid IDS

Modern switched networks have created a problem for intrusion detection operators. In a switched network, the NIC may be running in promiscuous mode, however the traffic may not be visible to the NIC. Some switches will not allow it at all, making the installation of a traditional network IDS difficult. Furthermore, high network speeds mean that many of the packets could be dropped by a NIDS. A solution has arisen in the form of Hybrid IDSs, which takes delegation of an IDS to a host one stage further, combining Network Node IDS and Host IDS in a single package. Recently the Hybrid IDS has been used to block attacks at the host and application level and vendor marketeers have called these **Host Intrusion Prevention Systems**.

Some vendors refer to any IDS that fills more than one role as being a Hybrid IDS. However, I feel this is more out of marketing greed than genuine honesty. The term "Hybrid IDS" was flavor of the month circa mid-2000 and many vendors wanted to jump on the bandwagon.

Network IDS (NIDS)

This monitors all network traffic passing on the segment where the sensor is installed, reacting to suspicious anomaly or signature-based activity. Traditionally these were promiscuous packet sniffers with IDS filters, though these days they have to be far more

intelligent, decoding protocols and maintaining state, etc. They analyze every packet for attack signatures, though under a heavy network load many will start to drop packets.

Many Network IDS have the facility to respond to attacks, which was covered under [Automated/Active Response](#) in Part One of this document. Hype about how Network IDS had seen its day pass due to high speeds and switched networks raises its head every so often, but some Network IDS can cope with gigabit speeds with minimal dropped packets, and switched networks can be overcome with spanning ports or TAPs such as those supplied by Shomiti.

Examples of Network IDS include [SecureNet Pro](#) and [Snort](#).

Network Node IDS (NNIDS)

Switched and/or high-speed networks have brought with them a problem: some Network IDS are unreliable at high speeds and when loaded they can drop a high percentage of the network packets. Switched networks often prevent a network IDS from seeing passing packets promiscuously. Network Node IDSs delegate the network IDS function down to individual hosts, alleviating the problems of both high speeds and switching. Network Node IDSs as a term never really took off and as their functionality increases, **Host Intrusion Prevention Systems** may be a more appropriate term.

Personal Firewall

Also known as a Host Intrusion Prevention System, personal firewalls sit on individual systems and prevent unwanted connections, incoming or outgoing. While not infallible, they are very effective in protecting hosts from attack. Not to be confused with Network Node IDS.

Examples include [ZoneAlarm](#) and [Sygate](#).

Target-Based IDS

This is one of those ambiguous IDS terms, which means different things to different people. One definition may refer to them being File Integrity Checkers, while an alternative is a Network IDS that only looks for signatures of attacks to which the protected network may be vulnerable. The objective of the latter definition is to speed up the IDS by not looking for unnecessary attacks. Personally, I want to know about every attack regardless of its chance of success. My point of view is that because the term has

meant very different things, it's use should be avoided for a few years -- in effect, quarantining the term to avoid confusion.

Network Intrusion Prevention System / Inline IDS

A Network Intrusion Prevention System sits inline on the network blocking attacks, similar to a firewall except it reacts using IDS signatures to accept or deny access. Whilst they can replace an IDS, in my opinion they aren't mature enough to replace a firewall. There are advantages and disadvantages in using them; I would suggest reading the definition of [automated response](#).

Host Intrusion Prevention System

A Host Intrusion Prevention System resides on the network host protecting it from attack. These used to be known as personal firewalls but as their capabilities increased the HIPS term took hold. There is a lot of vendor marketing surrounding the term and different products have vastly different functions.

Attack/DDOS Mitigation Tool

A DDOS attack succeeds by consuming sufficient network bandwidth to prevent any legitimate traffic reaching it's destination. An attack mitigation tool resides as close to the Internet as possible where the network bandwidth is generally larger. The idea is to block the DDOS attack whilst there is still sufficient unconsumed bandwidth, so as to not have an effect on the legitimate traffic. The best place for the product is at the ISP and unless it can be installed where the bandwidth is greater than the border router it will accomplish very little.

Additional IDS Terminology

Intrusion Detection Working Group (IDWG)

The purpose of the Intrusion Detection Working Group is to define data formats and exchange procedures for sharing information of interest to intrusion detection systems and response systems, and to management systems which may need to interact with them. The Intrusion Detection Working Group will coordinate its efforts with other IETF working groups.

Islanding

Cutting the network off from the Internet, islanding is a drastic action and used almost a last resort. It is occasionally used by some organizations faced with a large virus outbreak, or even

when the threat of an attack is considered significant enough.

Load Balancing

An IDS load balancer will divide traffic amongst various Network IDS based on source/destination addresses or port. Load balancers are required where traffic load is too much for the IDS to cope with or, less likely, where IDS redundancy is required.

Low and Slow

An attacker wanting to hide their intent will probe a host over a long period of time, one port every now and then randomly. This is very hard to detect except through historical advanced analysis and some good database queries. The trick for the attacker is to know the default settings for the detection systems they may face. My advice to the defender is to reduce these settings or carry out advanced analysis on historic data using commercial tools such as Silent Runner or bespoke SQL queries on the underlying database. The former methods may generate more false positives.

OS Fingerprinting

Active

Traditionally tools such as [nmap](#) would be used by an attacker to scope a target prior to an attack; now they are used within IDS to ascertain the threat posed by an attack based on how susceptible a target is to the attack. For instance, a windows attack against a Linux box would be considered benign. Active fingerprinting is identifying the operating system of a remote host through stimulus - response. Different operating systems respond to certain packets in different ways, allowing the fingerprinter to identify not only the OS but often its patch level.

Passive

Passive fingerprinting relies on using information within packets without stimulating their creation. Passive fingerprinting is slower than active fingerprinting as it relies on packets being generated naturally. It maps the network over time and may be unaware of hosts that have not had traffic pass by the IDS. However, being passive it allows the IDS to be completely stealthy especially if the IDS uses Data In Nothing Out (DINO) taps.

Out Of Band (OOB)

On a network an IDS can be stealthy enough to avoid detection. However the management

interface is still visible and therefore a prime target for attack. The management interface will often be taken Out Of Band, connected to a completely separate network. VPN's could also be used, creating a vOOB, but this isn't truly OOB as it is still susceptible to the same network DOS attacks as In Band reporting.

Promiscuous

By default, the IDS network interface only sees information addressed to it -- this is termed non-promiscuous. By making the interface promiscuous, you can see all the network traffic on your segment regardless of the source or destination address. This is essential for a Network IDS, but by the same token can be used by packet sniffers to monitor your network traffic. Switches make this monitoring much harder, however many have a span or mirror port which can be configured to enable you to see all the traffic.

Routers

A router is a device for connecting sub-networks. They operate at the Transport and Network layers of the OSI 7-layer model. More basically, routers help to navigate network packets to their destinations. Many also have Access Control Lists (ACL's) that will allow you to filter out undesirable packets. Many routers can feed their logs into an IDS correlation tool, providing valuable information about blocked attempts to access a network.

Shunning

Shunning is the practice of configuring border devices to reject any packets from objectionable sources. Some networks have been known to shun IP addresses from a specific country. More often than not ISPs with a poor clean up rate will be shunned.

Signatures

At the heart of the IDS is the attack signature, this is what makes the IDS trigger on an event. Too basic and it triggers too often, creating false positives. Too long and it slows the IDS down. Some people see the number of signatures that an IDS supports as a benchmark of the IDS's quality. However, while one vendor has one signature covering many attacks another vendor may list the signatures separately, giving the impression (to some) that because it appears to include more signatures, it's a better IDS.

Network Grepping / Pattern Matching

Some signatures are based on Network Grepping, looking for a sequence of traffic that matches that within an attack. Pattern matching or grepping signatures are easier to

create but are prone to reporting false positives.

Protocol Decode/Analysis

Protocol Decode/Analysis signatures understand various protocols and look for attack signatures within them, whilst slower than grepping they are not so prone to false positives.

Heuristic

Heuristic signatures look for abnormal traffic and events, combining various low priority events into a significant event or even repeated ASCII or binary characters in a single query.

Anomaly/Behavioral Signatures

Statistical Anomaly

A statistical anomaly based IDS highlights deviation from the general rule by building a profile of the host or network activity over time. When an event occurs which is outside this profile the IDS will alarm. For example, this happens in a Host IDS when a user suddenly performs a highly privileged function when he/she hasn't done so previously. Or, in the case of a network IDS, a profile is built of the network traffic over time, as this traffic shouldn't vary significantly without good reason. The IDS will then alert when the traffic steps outside certain parameters. As well as fulfilling a valuable security function the information is often extremely valuable to network administrators.

Protocol anomaly

These signatures identify packets that aren't compliant with a protocol RFC. Unfortunately many vendors do not comply with an RFC, which creates false positives. Others will look at a protocol with more thought, looking for traffic that would not be expected within a protocol regardless of whether is it RFC compliant. Some have referred to this latter method as a behavioral IDS. In my honest opinion there is very little difference between an anomaly and a process that doesn't follow normal behavior. I cynically suspect some vendors may be attempting to gain a marketing edge through new terminology, or less cynically, that the process does differ slightly therefore a new term is applicable to draw attention to this.

Stealth Mode

Stealth interfaces allow an IDS to be invisible to other hosts on the network while still being able to detect attacks. They have no TCP/IP stack. Where a packet needs to be transmitted through the interface the packet would need to be crafted. IDS will usually also have a management interface through which it can be managed.

Taps

To tap, in IDS terms, is to monitor the data on a live connection without being seen, similar to when law enforcement organizations would tap into telephone lines to listen to conversations. These are now available for copper and fiber and from 10MB/s to Gigabit. Things to look out for are fail safe, i.e., if the tap fails will it still pass traffic and is the taps output full duplex. IDS taps are not to be confused with a network TAP (Terminal Access Point) where you plug a host into the network.

Tarpit

A Tarpit slows down scans and probes destined for unassigned IP addresses, some have called it a "sticky honeypot". One such example is [LaBrea](#) which takes over unused IP addresses on a network and creates "virtual machines" that answer to connection attempts. LaBrea answers those connection attempts in a way that causes the machine at the other end to get "stuck", sometimes for a very long time. Tarpits are particularly effective at [slowing down worms](#) and scans.

Tuning

Tuning an IDS is adapting the signature policy to the environment in order to reduce false positives and address the threat. With some experience, coarse tuning can be carried out prior to deployment though fine tuning has to be carried out whilst the IDS is functioning. Tuning is an ongoing maintenance task for any IDS owner, though it becomes less painful once the IDS has bedded in.

Visualization

Most IDS will output their alerts in the form of a log. However, correlating these events can be very difficult even to experienced analysts. Viz tools will represent related events pictorially; for instance probed ports over time will identify low and slow scans. Another good one is source hosts connecting to target hosts through triggered events, the end result in 3D is like a dandelion seed ball, but the results jump out cutting down on time to analyse.

Summary

IDS terminology continues to develop, and some terms have even changed since I started writing this article. I have tried, where possible to include the various meanings of the same term and hope I have covered them to your satisfaction. If further clarification is sought regarding the terms or if you wish to discuss or comment on the terms please do not hesitate to contact me.

Andy Cuff is a computer security consultant who specializes in Intrusion Detection. He is currently responsible for deploying and maintaining various IDSs on a Global network of over 300,000 hosts. During the last two decades he has experienced a variety of different security roles ranging from cryptography to TEMPEST and from bug sweeping to pen testing. In his spare time he maintains the vendor independent [Talisker Security Tools](#) website which offers salient details on every known network security device. Andy is a regular contributor to many security-related mailing lists.

[Privacy Statement](#)

Copyright 2006, SecurityFocus